07-25-00 A

Please type a plus sign (+) inside this box → +

PTO/SB/05 (4/98)
Approved for use through 09/30/2000. OMB 0651-0032
Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

UTILITY PATENT APPLICATION TRANSMITTAL

(Only for new nonprovisional applications under 37 C.F.R. § 1.53(b))

Attorney Docket No.	12680ROUS02U
First Inventor or Application Identifier	R. Bruce Wallace
Title	Applications for Networked Storage Systems
Express Mail Label No.	EE430633413US

APPLICATION ELEMENTS

See MPEP chapter 600 concerning utility patent application contents.

ADDRESS TO: Assistant Commissioner for Patents, Box Patent Application, Washington, DC 20231

1. [X] * Fee Transmittal Form *(e.g., PTO/SB/17)* *(Submit an original and a duplicate for fee processing)*
2. [X] Specification [Total Pages 14] *(preferred arrangement set forth below)*
   - Descriptive title of the Invention
   - Cross References to Related Applications
   - Statement Regarding Fed sponsored R & D
   - Reference to Microfiche Appendix
   - Background of the Invention
   - Brief Summary of the Invention
   - Brief Description of the Drawings *(if filed)*
   - Detailed Description
   - Claim(s)
   - Abstract of the Disclosure
3. [X] Drawing(s) *(35 U.S.C. 113)* [Total Sheets 2]
4. Oath or Declaration [Total Pages 4]
   a. [X] Newly executed (original or copy)
   b. [ ] Copy from a prior application (37 C.F.R. § 1.63(d)) *(for continuation/divisional with Box 16 completed)*
      i. [ ] DELETION OF INVENTOR(S) Signed statement attached deleting inventor(s) named in the prior application, see 37 C.F.R. §§ 1.63(d)(2) and 1.33(b).

NOTE FOR ITEMS 1 & 13 IN ORDER TO BE ENTITLED TO PAY SMALL ENTITY FEES, A SMALL ENTITY STATEMENT IS REQUIRED (37 C.F.R. § 1.27), EXCEPT IF ONE FILED IN A PRIOR APPLICATION IS RELIED UPON (37 C.F.R. § 1.28).

5. [ ] Microfiche Computer Program *(Appendix)*
6. Nucleotide and/or Amino Acid Sequence Submission *(if applicable, all necessary)*
   a. [ ] Computer Readable Copy
   b. [ ] Paper Copy (identical to computer copy)
   c. [ ] Statement verifying identity of above copies

ACCOMPANYING APPLICATION PARTS

7. [X] Assignment Papers (cover sheet & document(s))
8. [ ] 37 C.F.R.§3.73(b) Statement *(when there is an assignee)* [ ] Power of Attorney
9. [ ] English Translation Document *(if applicable)*
10. [ ] Information Disclosure Statement (IDS)/PTO-1449 [ ] Copies of IDS Citations
11. [ ] Preliminary Amendment
12. [X] Return Receipt Postcard (MPEP 503) *(Should be specifically itemized)*
13. [ ] * Small Entity Statement(s) *(PTO/SB/09-12)* [ ] Statement filed in prior application Status still proper and desired
14. [ ] Certified Copy of Priority Document(s) *(if foreign priority is claimed)*
15. [ ] Other: .....

16 **If a CONTINUING APPLICATION,** *check appropriate box, and supply the requisite information below and in a preliminary amendment:*

[ ] Continuation [ ] Divisional [ ] Continuation-in-part (CIP) of prior application No: _____/_____

Prior application information *Examiner* _____ *Group / Art Unit:* _____

For CONTINUATION or DIVISIONAL APPS only: The entire disclosure of the prior application, from which an oath or declaration is supplied under Box 4b, is considered a part of the disclosure of the accompanying continuation or divisional application and is hereby incorporated by reference. The incorporation can only be relied upon when a portion has been inadvertently omitted from the submitted application parts.

17. CORRESPONDENCE ADDRESS

[ ] *Customer Number or Bar Code Label* *(Insert Customer No. or Attach bar code label here)* or [X] *Correspondence address below*

Name	Holmes W. Anderson				
	Nortel Networks Limited				
Address	600 Technology Park Drive				
	Mail Stop E65-60-403				
City	Billerica	State	MA	Zip Code	01821-5501
Country		Telephone	001-978-288-6997	Fax	001-978-288-6543

Name (Print/Type)	Holmes W. Anderson	Registration No. (Attorney/Agent)	37,272
Signature	Holmes W. Anderson	Date	24 July 2000

Burden Hour Statement: This form is estimated to take 0.2 hours to complete. Time will vary depending upon the needs of the individual case. Any comments on the amount of time you are required to complete this form should be sent to the Chief Information Officer, Patent and Trademark Office, Washington, DC 20231. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Assistant Commissioner for Patents, Box Patent Application, Washington, DC 20231

12680ROUS02U

Express Mail Label No. EE430633413US
Date of Mailing: July 24, 2000

Inventors: Bruce Wallace
Bruce Townsend
Tom Chmara
Ryan Stark
John Bell

APPLICATIONS FOR NETWORKED STORAGE SYSTEMS

FIELD OF THE INVENTION

The present invention is generally related to network communications, and more particularly to applications for networked storage systems.

CROSS-REFERENCE TO RELATED APPLICATIONS

A claim of priority is made to U.S. Provisional Patent Application Serial No. 60/208,472, entitled Data Access in a Network Environment, filed on June 2, 2000.

BACKGROUND OF THE INVENTION

Techniques that enable multiple computers to access multiple data storage devices in a network are known. For example, personal computers may access a Redundant Array of Independent Disks (RAID) through a server. The personal computers may communicate with the server via a Local Area Network (LAN) employing Ethernet, Token Ring or FDDI, while the server communicates with the storage devices associated with the RAID in accordance with a non-network protocol such as Small Computer System Interface (SCSI). However, such known configurations have limitations. For example, the number of data storage devices that can be addressed via each SCSI interface is limited to only 4, 8 or 16, depending on the version of SCSI that is implemented. Further, the physical characteristics of the SCSI cabling that connects the data storage devices to the server impose limitations on the geographic placement of the data storage devices. In particular, the server and the data storage devices must be located in relatively close physical proximity because the maximum practical length of the SCSI cabling is relatively short. Such proximity

results in vulnerability to catastrophic data loss due to events such as fire. While techniques for communicating SCSI transmissions over larger distances via Internet Protocol (IP) and Fiber Channel (FC) have been proposed, some of the problems remain.

SUMMARY OF THE INVENTION

In accordance with the present invention, at least one network device that moves storage blocks based upon a session layer (Open Systems Interconnection ("OSI") Layer 5) or higher protocol set is employed to facilitate communication associated with storage between at least one host and at least one target storage device. The network device may be a switch. Further, the switch may be situated within the network, rather than at the network edge. Transactions related to data storage access may be expressed in a non-network protocol and transmitted via a network protocol data unit across at least a portion of the network by employing techniques such as encapsulation.

The present invention enables network data storage implementations with enhanced performance and reduced likelihood of catastrophic data loss. Performance is enhanced because a relatively large number of data storage devices may be addressed, thereby enabling N-way striping and mirroring. The likelihood of catastrophic data loss is reduced because the physical distance between individual data storage devices can be increased. For example, data stored in a first data storage device may be mirrored by a second data storage device that is located at a sufficiently great physical distance from the first data storage device to mitigate the possibility of contemporaneous destruction of both data storage devices by an event such as a fire. Flexibility in the choice of physical location for data storage devices combined with N-way mirroring enables load balancing and storage of selected data on a faster communication path relative to selected network devices, such as devices that most frequently access the data, without precluding data access by other network devices. Further, network scalability is facilitated by simplifying network connections. Hence, many devices may share a file store and data storage devices may be easily added to the network.

In another embodiment of the present invention, a READ operation is facilitated by transmitting READ metadata to the filesystem associated with the client.

The READ metadata describes a procedure for accessing data stored on at least one networked storage device, and may be maintained by a filesystem associated with a content owner device. When a client transmits a READ transaction to the networked storage device, READ metadata is transmitted from the networked storage device to the client device. The transmitted READ metadata is selected to be sufficient to enable the client device to obtain the data associated with the READ operation from the networked storage device without communication with the content owner device. The client employs the metadata with a "read only" read utility to obtain the data. Since a "local copy" of the metadata is employed and the content owner device is not included in the read transaction, performance is improved. The content owner device may issue an update by providing new metadata to the networked storage device.

BRIEF DESCRIPTION OF THE DRAWING

. The invention will be more fully understood from the following Detailed Description and Drawing, of which:

Fig. 1 is a block diagram of a portion of a network that illustrates use of a storage services protocol set to facilitate network storage; and

Fig. 2 is a block diagram of a portion of a network that illustrates a READ operation in a network where the data is owned by a content owner device.

DETAILED DESCRIPTION

Fig. 1 is a block diagram of a portion of a network. Client devices 8, 9, 10, 11, 12, 13, 14 are connected with a plurality of data storage units 16, 18, 20, 22, 24, 26 via hosts 33, 34, 35 and a switch 50 capable of moving storage blocks based upon a session layer or higher ("storage services") protocol set. Each device in the network is associated with a network address, such as an Internet Protocol (IP) address. Each data storage unit includes an interface adapter 38 and at least one physical storage device 40 such as a disk.

The storage services protocol set is employed to facilitate transmission of data units associated with storage operations between the clients and the data storage units. Transactions between the clients and the hosts may be represented in network protocol data units 42, such as packets, frames or cells. In the illustrated embodiment,

communication of transactions between the clients and the hosts is in accordance with Transport Control Protocol (TCP) or Unacknowledged Datagram Protocol (UDP) over IP over Ethernet. Corresponding communications between the hosts, switch 50 and data storage units 16-26 employ modified representations of the network protocol data units wherein SCSI, IDE or similar command sets are employed by the storage services protocol set. In particular, corresponding non-network protocol transactions are generated and encapsulated in network protocol data units 44 to provide, for example, SCSI or IDE over TCP or UDP over IP over Ethernet. Those skilled in the art will recognize that other protocols may be utilized.

The hosts employ metadata to generate the encapsulating data units. For example, in a READ operation initiated by the client 10, a representation of at least one block or file to be retrieved is transmitted from the client to the host 34 via data unit 42. Following receipt of the data unit 42, the host 34 determines where data that corresponds to the READ operation is located at the storage devices by employing metadata that is maintained on a filesystem 46 associated with the host 34. The metadata describes the directory hierarchy, file attributes, permissions, and file layout in terms of a block address chain. The host may first employ the metadata to determine whether to allow the client access to the data. Techniques for determining whether to allow access are well known in the art. If access to the data is to be allowed, the host employs the metadata to generate a corresponding non-network protocol READ transaction that is encapsulated in at least one network protocol data unit 44. For example, a SCSI transaction and a header portion may be generated and combined to form the new protocol data unit 44, where the SCSI transaction is placed in the payload portion of the data unit. The new network protocol data unit 44 is then transmitted to the switch 50, which directs the data unit 44 to the appropriate data storage unit 26 based on predetermined criteria, such as IP destination address and storage services protocol set information.

The data storage units operate upon the encapsulated transactions following receipt thereof. For example, the READ transaction is executed by retrieving the data that is specified by the encapsulated SCSI transaction and returning that data to the host 34. Following receipt of the network protocol data unit 44 by the designated data storage unit 26, the encapsulated transaction is extracted from the data unit. In particular, the interface adapter 38 associated with the data storage unit 26 is employed to strip the header portion of the data unit from the payload portion. The IP

source address and other information from the header portion may be retained to facilitate transmission of data from the data storage unit to the host. The remaining non-network protocol transaction is employed to retrieve data in accordance with techniques known in the art. The interface adapter 38 is employed to encapsulate the retrieved data in at least one network protocol data unit 46 that designates the IP address of the host in the destination field. The network protocol data unit 44 is transmitted to the host 34 via the switch 50 based on predetermined criteria, such as IP destination address and storage services protocol set information. Following receipt of the data unit 44 that contains the data, the host generates at least one corresponding network protocol data unit 48 that indicates the IP address of the client in the destination address field. The client 10 extracts the data from the data unit 48.

The predetermined criteria for directing the data units may include employing the storage services protocol set at the switch or elsewhere to facilitate implementation of the data storage units as a Redundant Array of Independent Disks (RAID). If data is mirrored, striped or otherwise located on a plurality of the data storage units, the data units associated with the transactions are replicated at the switch based on storage services protocol set information. Reliable unicast or reliable multicast may be employed by the switch to transmit the data units between the hosts to the data storage units.

The storage services protocol set may also be employed to enhance performance by balancing the data access load and testing for data integrity. In the case of a READ operation where mirroring is implemented, load data may be employed by the switch to determine the data access load on each data storage unit, and to balance the load on the data storage units by transmitting the data unit associated with the transaction only to a selected data storage unit with relatively less load. Alternatively, the data unit may be transmitted to each mirrored data storage unit in order to obtain data from the fastest responding mirrored data storage unit. Alternatively, data returned from a plurality of mirrored data storage units may be compared by the switch to provide greater assurance of data integrity.

In order to provide protection against catastrophic data loss, the mirrored data storage units may be separated by relatively large physical distances. In mirroring, multiple identical copies of data are maintained on multiple disks, typically one copy per disk where each data storage unit includes at least one disk. The disks associated with mirrored pairs are separated by a sufficiently large distance to protect against

contemporaneous destruction by a catastrophic event such as fire. Further, performance may be enhanced by implementing an N-way mirrored RAID, where N is limited only by the number of available IP addresses. Individual mirrored disks associated with the N-way mirrored RAID are disposed along a respective fast communication path relative to areas of greater network activity. For example, a mirrored disk may be located near each one of a group of urban areas.

Referring now to Fig. 2, data access may also be enhanced by providing metadata to the clients. In the illustrated network portion, hosts are not employed and data ownership is asserted by a content owner device 56. The content owner 56 retains control of READ and WRITE privileges for at least some of the data stored in the data storage units 16-26. However, metadata from filesystem 62 is provided to the data storage units by the content owner device via at least one data unit 60. To obtain owned data, the client 10 initially transmits a READ transaction to at least one of the data storage units. The READ transaction employs at least one network protocol data unit 58 that designates at least one file or block of data that is to be transmitted to the client. The data unit 58 is transmitted from the client to the Layer 3 router 36, and from the Layer 3 router to the designated data storage units. The receiving data storage units process the READ transaction to determine whether to grant the client access to the data. Techniques for determining whether to allow access are well known in the art. If access to the data is to be allowed, the data storage unit selects READ metadata that describes the location of the data associated with the READ transaction in terms of a block address chain. The selected READ metadata is transmitted from the data storage unit to the client via the Layer 3 router in at least one network protocol data unit 70.

The transmitted READ metadata is employed by the filesystem 64 of the client 10 to directly access the associated data from the data storage units. A new READ transaction is generated by the client and transmitted to the data storage units via at least one network protocol data unit 72. In particular, a non-network protocol READ transaction is placed in a payload portion of at least one protocol data unit, and a header portion is generated for the protocol data unit. The header portion of the protocol data unit 72 includes a source address field that contains the IP source address that is associated with the client 10, and a destination address field that contains the IP destination address that is associated with the appropriate data storage unit 22. The data unit is then transmitted to the Layer 3 router, which routes the data

unit to the data storage unit designated by the IP destination address. Following receipt of the network protocol data unit by the data storage unit, the encapsulated non-network protocol READ transaction is extracted from the data unit 72. In particular, the interface adapter 38 associated with the data storage unit 22 is employed to strip the header portion of the data unit from the payload. The IP source address and other information from the header portion may be retained to facilitate transmission of data from the data storage unit to the client. The READ transaction is executed by retrieving the data that is specified by the READ transaction and returning that data to the client. As described above, the interface adapter is employed to place the retrieved data in at least one network protocol data unit that designates the IP address of the client in the destination field. The network protocol data units are then transmitted to the client via the Layer 3 router. A host could alternatively be employed to mitigate the client's computational overhead associated with encapsulating non-network protocol transactions. The content owner device transmits updated metadata to the data storage units when the storage content changes.

The described technique is compatible with RAID implementations and mirrored sites in an Internet environment. In the case of a READ operation where mirroring is implemented, the READ metadata that is transmitted from the content owner device to the data storage units and from the data storage units to the client may each be selected to achieve one or more performance objectives. For example, load data may be employed to determine the data access load on each data storage unit, and balance the load on the data storage units by transmitting READ metadata that corresponds to a selected data storage unit with relatively less load. Alternatively, READ metadata corresponding to a data storage unit that is located along a relatively faster communication path to the client relative to other data storage units may be transmitted to the client. The READ metadata may alternatively correspond to a plurality of data storage units in order to obtain data from the fastest responding data storage unit.

Variations, modifications and different applications of the invention taught in the Detailed Description and Drawing will be apparent to those skilled in the art. For example, it will be understood by those skilled in the art that the described storage services protocol set based processing functions may be performed by a device other than the illustrated switch. Accordingly, it should be understood that other

embodiments of the invention might be practiced without departing from the spirit and scope of the invention as defined by the appended claims.

CLAIMS

What is claimed is:

1. A method for facilitating operations related to data storage between a first device and at least one data storage unit in a computer network, comprising the steps of:

causing generation of a non-network protocol transaction;

representing the non-network protocol transaction in at least one network protocol data unit;

processing at least one network protocol data unit based on a storage services protocol set to facilitate transmission of the data unit;

extracting the non-network protocol transaction from the network protocol data unit; and

operating upon the non-network protocol transaction.

2. The method of claim 1 wherein said processing step includes the further step of balancing loads associated with selected read transactions.

3. The method of claim 1 wherein said processing step includes the further step of duplicating data units associated with selected write transactions to achieve mirroring.

4. The method of claim 1 wherein said processing step includes the further step of duplicating data units associated with selected transactions to achieve N-way mirroring.

5. The method of claim 1 wherein said processing step includes the further step of duplicating selected metadata.

6. The method of claim 1 wherein said processing step includes the further step of ensuring right to access based on originator.

7. The method of claim 1 wherein said processing step includes the further step of blocking access to selected destinations.

8. The method of claim 1 wherein said processing step includes the further step of monitoring and logging access.

9. The method of claim 8 wherein said processing step includes the further step of employing results from access monitoring and logging to detect unauthorized intrusion.

10. Apparatus that facilitates operations related to data storage between a first device and at least one data storage unit in a computer network, comprising:

a filesystem that indicates location of data stored on at least one data storage unit, and

circuitry that processes network protocol data units associated with the operations based on storage services protocol set information to facilitate transmission of the data unit.

11. The apparatus of claim 10 wherein said circuitry balances loads associated with selected read transactions.

12. The apparatus of claim 10 wherein said circuitry duplicates data units associated with selected write transactions to achieve mirroring.

13. The apparatus of claim 10 wherein said circuitry duplicates data units associated with selected transactions to achieve N-way mirroring.

14. The apparatus of claim 10 wherein said circuitry duplicates selected metadata.

15. The apparatus of claim 10 wherein said circuitry facilitates ensuring right to access based on originator.

16. The apparatus of claim 10 wherein said circuitry blocks access to selected destinations.

17. The apparatus of claim 10 wherein said circuitry monitors and logs access.

18. The apparatus of claim 17 wherein said circuitry employs results from access monitoring and logging to detect unauthorized intrusion.

19. The apparatus of claim 10 wherein said circuitry includes a host device.

20. The apparatus of claim 10 wherein said circuitry includes a switch.

21. A method for facilitating a READ operation that is initiated by a client transmitting a READ transaction to at least one data storage device for data that is owned by a content owner device, comprising the steps of:

in response to said READ transaction, selecting READ metadata previously transmitted from a filesystem that is associated with the content owner to the data storage device; and

transmitting the READ metadata from the data storage device to the client via at least one network protocol data unit.

22. The method of claim 21 including the further step of determining whether to grant the client access to the data.

23. The method of claim 22 including the further step of employing the transmitted READ metadata with a filesystem that is associated with the client to access the data from the at least one data storage unit.

24. The method of claim 23 wherein said step of employing the READ metadata includes the further step of generating a new READ transaction and transmitting the new READ transaction to the data storage units via at least one network protocol data unit.

25. The method of claim 24 wherein said step of generating a new READ transaction includes the further step of placing a non-network protocol READ transaction in a payload portion of at least one protocol data unit and generating a header portion for the protocol data unit.

26. Apparatus that facilitates a READ operation that is initiated by a client that transmits a READ transaction, comprising:

at least one data storage unit that includes metadata that indicates location of data at said at least one data storage unit, said data storage unit including circuitry that is operative in response to receipt of the READ transaction from the client to:

select READ metadata that corresponds to the READ operation from the content owner, and

transmit the selected READ metadata to the client via at least one network protocol data unit,

whereby the client might employ said READ metadata to obtain data directly from said at least one data storage unit.

27. The apparatus of claim 26 wherein said data storage unit includes circuitry that determines whether to grant the client access to the data based upon predefined criteria.

28. The apparatus of claim 27 further including circuitry associated with the client that is operative to generate a non-network protocol READ transaction based at least in part upon said transmitted READ metadata, and encapsulate said non-network protocol READ transaction in at least one network protocol data unit having a destination address field.

29. The apparatus of claim 28 further including a Layer 3 router that routes the at least one network protocol data unit to the data storage unit designated by the destination address.

30. The apparatus of claim 29 wherein said designated data storage unit includes an interface adapter operative to extract the encapsulated non-network protocol READ transaction from the network protocol data unit.

31. The apparatus of claim 30 wherein said designated data storage unit is operative in response to said non-network protocol READ transaction to retrieve the data that is specified by the READ transaction, and wherein said interface adapter is

further operative to place the retrieved data in at least one network protocol data unit that designates the address of the client in a destination field.

ABSTRACT

At least one network device that moves storage blocks based upon a session layer or higher protocol set is employed to facilitate communication between at least one host and at least one target storage device. The at least one network device may include a server and a switch. Transactions related to data storage access are expressed in a non-network protocol and transmitted via a network protocol data unit across at least a portion of the network by employing encapsulation. An adapter card associated with the data storage unit is employed to extract the non-network protocol transaction. Multiple, mirrored data storage units that are separated by relatively large physical distance may be employed to reduce the likelihood of catastrophic data loss due to an event such as a fire. Further, READ transactions may be facilitated by preloading data storage units with READ metadata for subsequent transmission to a client device to enable expedited access to data stored on the data storage units.

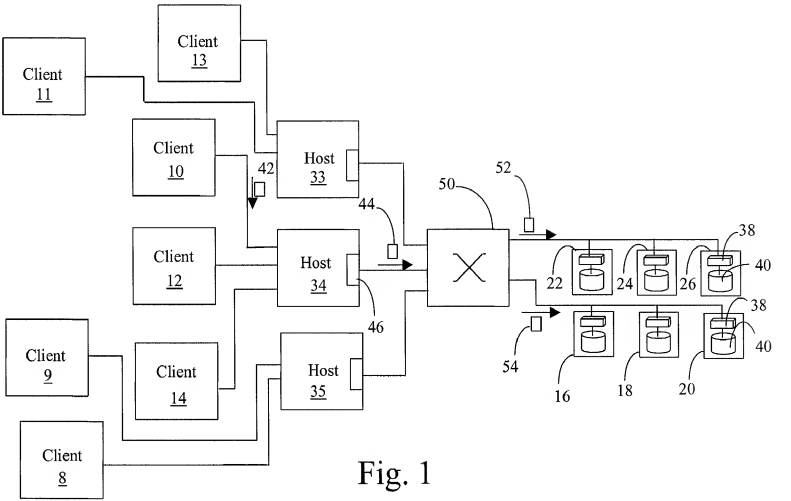

Fig. 1

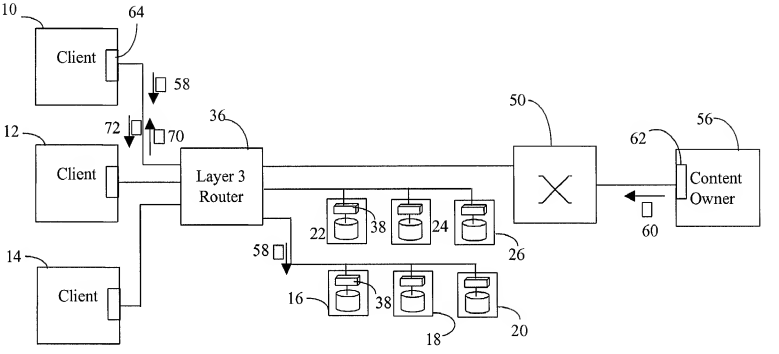

Fig. 2

Express Mail Label No. EE430633413US
Date of Mailing July 24, 2000

Please type a plus sign (+) inside this box → +

PTO/SB/01 (12-97)
Approved for use through 9/30/00. OMB 0651-0032
Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

DECLARATION FOR UTILITY OR DESIGN PATENT APPLICATION (37 CFR 1.63)

☒ Declaration Submitted with Initial Filing OR ☐ Declaration Submitted after Initial Filing (surcharge (37 CFR 1.16 (e)) required)

Attorney Docket Number	12680ROUS02U
First Named Inventor	Wallace, R. Bruce
COMPLETE IF KNOWN	
Application Number	
Filing Date	
Group Art Unit	
Examiner Name	

As a below named inventor, I hereby declare that:

My residence, post office address, and citizenship are as stated below next to my name

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled:

Applications for Networked Storage Systems

(Title of the Invention)

the specification of which
☒ is attached hereto
OR
☐ was filed on (MM/DD/YYYY) ____ as United States Application Number or PCT International Application Number ____ and was amended on (MM/DD/YYYY) ____ (if applicable).

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment specifically referred to above.

I acknowledge the duty to disclose information which is material to patentability as defined in 37 CFR 1.56.

I hereby claim foreign priority benefits under 35 U.S.C. 119(a)-(d) or 365(b) of any foreign application(s) for patent or inventor's certificate, or 365(a) of any PCT international application which designated at least one country other than the United States of America, listed below and have also identified below, by checking the box, any foreign application for patent or inventor's certificate, or of any PCT international application having a filing date before that of the application on which priority is claimed.

Prior Foreign Application Number(s)	Country	Foreign Filing Date (MM/DD/YYYY)	Priority Not Claimed	Certified Copy Attached? YES	NO
			☐	☐	☐
			☐	☐	☐
			☐	☐	☐
			☐	☐	☐

☐ Additional foreign application numbers are listed on a supplemental priority data sheet PTO/SB/02B attached hereto.

I hereby claim the benefit under 35 U.S.C. 119(e) of any United States provisional application(s) listed below.

Application Number(s)	Filing Date (MM/DD/YYYY)	
60/208,472	06/02/2000	☐ Additional provisional application numbers are listed on a supplemental priority data sheet PTO/SB/02B attached hereto.

Please type a plus sign (+) inside this box → [+]

PTO/SB/01 (12-97)
Approved for use through 9/30/00. OMB 0651-0032
Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

DECLARATION — Utility or Design Patent Application

I hereby claim the benefit under 35 U.S.C. 120 of any United States application(s), or 365(c) of any PCT international application designating the United States of America, listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States or PCT international application in the manner provided by the first paragraph of 35 U.S.C. 112, I acknowledge the duty to disclose information which is material to patentability as defined in 37 CFR 1.56 which became available between the filing date of the prior application and the national or PCT international filing date of this application.

U.S. Parent Application or PCT Parent Number	Parent Filing Date (MM/DD/YYYY)	Parent Patent Number (if applicable)

☐ Additional U.S. or PCT international application numbers are listed on a supplemental priority data sheet PTO/SB/02C attached hereto.

As a named inventor, I hereby appoint the following registered practioner(s) to prosecute this application and to transact all business in the Patent and Trademark Office connected therewith: ☐ Customer Number OR ☒ Registered practitioner(s) name/registration number listed below → Place Customer Number Bar Code Label here

Name	Registration Number	Name	Registration Number
Holmes W. Anderson	37,272	Lindsay G. McGuinness	38,549
Christopher J. Cianciolo	42,417	Mary M. Steubing	37,946
John C. Gorecki	38,471		

☐ Additional registered practitioner(s) named on supplemental Registered Practitioner Information sheet PTO/SB/02C attached hereto.

Direct all correspondence to: ☐ Customer Number or Bar Code Label [ ] OR ☒ Correspondence address below

Name	Holmes W. Anderson, Nortel Networks Limited				
Address	600 Technology Park Drive				
Address	Mail Stop E65-60-403				
City	Billerica	State	MA	ZIP	01821-5501
Country	U.S.A.	Telephone	001-978-288-6997	Fax	001-978-288-6543

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under 18 U.S.C. 1001 and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Name of Sole or First Inventor: ☐ A petition has been filed for this unsigned inventor

Given Name (first and middle [if any])	Family Name or Surname
R. Bruce	Wallace

Inventor's Signature	[signature]					Date	JULY 19, 2000
Residence: City	Ashton, Ontario	State		Country	Canada	Citizenship	Canada
Post Office Address	18 Carlisle Circle, RR #4						
Post Office Address							
City	Ashton, Ontario	State		ZIP	K0A1B0	Country	Canada

☒ Additional inventors are being named on the 2 supplemental Additional Inventor(s) sheet(s) PTO/SB/02A attached hereto

Please type a plus sign (+) inside this box → +

PTO/SB/02A (3-97)
Approved for use through 9/30/98. OMB 0651-0032
Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

DECLARATION

ADDITIONAL INVENTOR(S) Supplemental Sheet
Page 1 of 2

Name of Additional Joint Inventor, if any:	☐ A petition has been filed for this unsigned inventor					
Given Name (first and middle [if any])			Family Name or Surname			
Bruce L.			Townsend			
Inventor's Signature				Date	July 19, 2000	
Residence: City	Stittsville, Ontario	State		Country Canada	Citizenship	Canada
Post Office Address	92 Hobin Street					
Post Office Address						
City	Stittsville, Ontario	State		ZIP K2S1G8	Country	Canada

Name of Additional Joint Inventor, if any:	☐ A petition has been filed for this unsigned inventor					
Given Name (first and middle [if any])			Family Name or Surname			
Thomas P.			Chmara			
Inventor's Signature				Date	19 Sept 2000	
Residence: City	Richmond, Ontario	State		Country Canada	Citizenship	Canada
Post Office Address	108 Fowler Street P.O. Box 952					
Post Office Address						
City	Richmond, Ontario	State		ZIP K0A2Z0	Country	Canada

Name of Additional Joint Inventor, if any:	☐ A petition has been filed for this unsigned inventor					
Given Name (first and middle [if any])			Family Name or Surname			
Ryan			Stark			
Inventor's Signature				Date	20/07/00	
Residence: City	Ottawa, Ontario	State		Country Canada	Citizenship	Canada
Post Office Address	134 Kenilworth Street					
Post Office Address						
City	Ottawa, Ontario	State		ZIP K1Y3Z5	Country	Canada

Burden Hour Statement: This form is estimated to take 0.4 hours to complete. Time will vary depending upon the needs of the individual case. Any comments on the amount of time you are required to complete this form should be sent to the Chief Information Officer, Patent and Trademark Office, Washington, DC 20231. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Assistant Commissioner for Patents, Washington, DC 20231

Please type a plus sign (+) inside this box → +

PTO/SB/02A (3-97)
Approved for use through 9/30/98. OMB 0651-0032
Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

DECLARATION

ADDITIONAL INVENTOR(S)
Supplemental Sheet
Page 2 of 2

Name of Additional Joint Inventor, if any:	☐ A petition has been filed for this unsigned inventor	
Given Name (first and middle [if any])	Family Name or Surname	
John H.	Bell	
Inventor's Signature		Date 19/7/00
Residence: City Ottawa, Ontario	State	Country Canada
Citizenship Canada		
Post Office Address 95 Bradford Street		
Post Office Address		
City Ottawa, Ontario	State	ZIP K2B5Y7
Country Canada		

Name of Additional Joint Inventor, if any:	☐ A petition has been filed for this unsigned inventor	
Given Name (first and middle [if any])	Family Name or Surname	
Inventor's Signature		Date
Residence: City	State	Country
Citizenship		
Post Office Address		
Post Office Address		
City	State	ZIP
Country		

Name of Additional Joint Inventor, if any:	☐ A petition has been filed for this unsigned inventor	
Given Name (first and middle [if any])	Family Name or Surname	
Inventor's Signature		Date
Residence: City	State	Country
Citizenship		
Post Office Address		
Post Office Address		
City	State	ZIP
Country		